



Exceptional service in the national interest



Sandia National Laboratories

LDRD IMPACTS ON THE SANDIA NUCLEAR WEAPONS PROGRAM

FROM NANOCAGES TO SUPERCOMPUTERS





From the Chief Technology Officer

The LDRD Program Innovates for Sandia's Nuclear Weapons Mission

The face of Sandia's Nuclear Weapons (NW) program is undergoing dramatic change as our Nation embarks on a period of unprecedented refurbishment of the stockpile, while simultaneously ensuring its safety, security, and reliability. While the bulk of the NW program must, of necessity, focus on meeting near-term commitments to the program, the Laboratory must simultaneously be tapping the creative juices of its staff to innovate options to meet future Nuclear Weapons needs.

While portions of the NW program are devoted to exploring future technologies and options for the stockpile, Sandia's Laboratory Directed Research and Development (LDRD) program also has a key role to play. LDRD projects enable Sandia scientists and engineers to both execute Sandia's broad national security mission and advance the frontiers of science and engineering as they seek understanding that leads to improvements in extant processes and materials and innovative approaches to achieving key functions needed for the NW stockpile. Since its inception by Congress in 1992, the Sandia LDRD program has been a major contributor to scientific understanding and technological innovations that help render the US nuclear stockpile a more secure and reliable deterrent, as well as others that contribute to the nonproliferation of special nuclear materials (SNM).

This brochure is designed to give the reader a flavor for some specific contributions of Sandia's LDRD program to the US nuclear weapons program, and it samples three general types of impact. First, there are LDRD projects that—through a series of accomplishments over time—have engendered impact that has already been widely recognized. Second, there are projects whose impact is more recent and, to some extent, in progress. And finally, there are LDRD initiatives whose impact is expected to help foster NW program innovations that are yet to be realized. This record is indicative of a program that is continually seeking opportunities to facilitate positive, creative evolution within one of Sandia's most important national security mission spaces.

In all cases, technical impact on the NW program is the result of fundamental science and engineering research—from materials science to chemistry to nuclear physics to microelectronics to high performance computing. This is the hallmark of LDRD-funded projects: fundamental scientific and engineering research leads to enhanced understanding and ability to manipulate physical systems, in turn leading to the application of such novel capabilities to crucial challenges in weapons science and technology. Delivering on Sandia's NW mission while advancing the frontiers of science, technology, and engineering—"innovation" and "nuclear weapons" are inextricably entwined. NNSA and the Nation substantially benefit from these LDRD-initiated innovations that help ensure improved stockpile performance, and therefore, enhanced national security.

J. Stephen Rottler

CTO and Vice President

Science, Technology, and Engineering

Research Foundations

jsrottl@sandia.gov

505-844-3882



Sandia National Laboratories

LDRD IMPACTS ON THE SANDIA
NUCLEAR WEAPONS PROGRAM

FROM NANOCAGES TO SUPERCOMPUTERS

CONTENTS

Developing High-Impact Codes for Massively Parallel Computing	4
Designing Clever Molecular-capture Cages for Nonproliferation and Sensing Critical Environments	8
Radiation Hardened Electronics for Critical Systems	12
Ingenious Lubrication Schemes for Preventing Wear in Critical Electromechanical Components	16
Encapsulating and Protecting Delicate Electronics	20
More-reliable Methods of Verification for FPGA-based Digital Systems	24

Developing High-Impact Codes for Massively Parallel Computing

Software development impacting numerous science and engineering arenas



Figure 1. Physical array of racks of processors composing a massively parallel (MP) computer, in this case, Sandia's experimental CPlant machine.

“It can be safely said that Trilinos would probably not exist if not for that LDRD.”

With the advent of high-performance computing (HPC) in the form of massively parallel (MP) supercomputers, an entire new generation of software engineering has been initiated—much of it LDRD supported—in order to facilitate the application of this computational power to the broadest possible array of national security challenges. These include weapons simulation and other weapons-related codes, which ensure the full functionality of the stockpile without actual testing.

THE SOLVER LIBRARY APPROACH

MP machines (fig. 1) are greatly expanded versions of what the PC world knows as multicore processing, that is, hardware that can process more than one data stream simultaneously. In MP supercomputers, the scale is huge: thousands of independent hardware processors, some or all multicore, crunching data simultaneously. Software programming for these machines quickly became a time-intensive enterprise, until the realization that there were certain algorithms in such software that tended to appear in programs across the sciences—from life science to electrical engineering to weapons physics.

Sandians initially took on the challenge of building libraries of such algorithms in an LDRD project known as “Hybrid Sparse-Dense Incomplete Factorization Preconditioners.” This software enabled the whole concept of modular, reusable library components as independent but interoperable packages available to programmers for use in diverse science and engineering applications. The original three packages in this

LDRD project ultimately spawned Trilinos, in 2004 — the first software library of algorithms for HPC science and engineering applications.

Over the decade from 1997–2006, this novel approach to software for MP computing garnered three R&D100 Awards: Aztec (1997), Trilinos (2004) and Compute Process Allocator (CPA) (2006). All three discovered ingenious methods to enhance the utility of MP computing without the significant investment costs of adding hardware. Instead, they provided programmers with the means to more efficiently utilize the extant processing power.

These Sandia approaches to improving the power and reliability of MP computing took two complementary paths, both originating in LDRD-funded projects. The first path facilitates the task of software developers in writing code to solve specific problems, the second optimizes the way that such software routines and subroutines actually run on MP machines from the standpoint of hardware allocation. MP computers have their drawbacks, and unreliability can be one of them. For example, estimates are that in a machine that contains thousands of individual processors, one of these processors may fail as frequently as every 15 minutes.

Emerging from those 1990s LDRD projects, Aztec constituted a parallel solver, that is, a software package designed to assist in finding the values for unknowns in a system of equations representing some large-scale scientific or engineering challenge, and running on an MP machine. In addition to climate modeling, examples of complex, real-world problems of this type would be transport of water-dissolved toxins through the environment, and the physical description of electrical instabilities and defibrillation of a human heart that has lost its intrinsic pacemaker rhythm. Aztec was originally designed as a solver for the application, MPSalsa, which approached the problem of simulating the dynamics of fluid flow, including applications such as the dynamics of flows occurring within a chemical vapor deposition reactor, (crucial to the process by which, for example, thin semiconductor films are created for computer chips). The solver package was subsequently generalized to render it useful in a broader variety of applications.

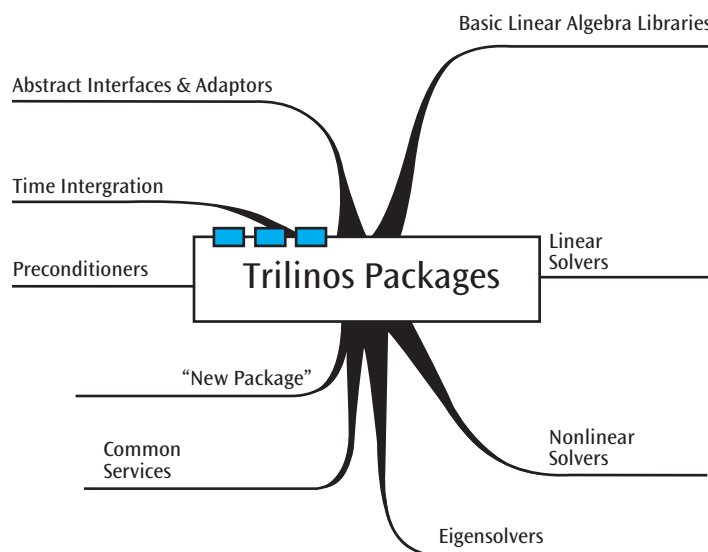


Figure 2. Overview of some of the software packages that make up Trilinos.

TRILINOS AND ITS OFFSHOOTS

This trend toward a software library of solvers was continued, refined, and rendered far more powerful in Trilinos—a software library whose components or “packages” can be adapted by way of application programming interfaces (APIs, quite common in the C++ programming language) to running as solvers for a variety of specific software applications. “In the early 1970s, almost no libraries were used in applications, but over time, we’ve progressively come to embrace them,” according to Trilinos principal developer, Mike Heroux.

The general notion is that in writing software for a specific problem to be modeled on an MP machine, a software developer will necessarily need to craft code that is germane to the specific issue under examination — climate modeling or weapons simulations, for example. However, there are also more generic aspects of that software construction connected with solving groups of differential equations whose solvers can be found in one or more of Trilinos’ packages. The developer who can make use of those generic, broadly applicable Trilinos packages (fig. 2)—plugging them into the situation-specific code—has greatly facilitated his/her task. By analogy, your computer might accept a mouse manufactured by several different companies, as long as they all possess the same standard USB-plug interface. In the case of Trilinos, a software package modeling climate and software modeling weapons dynamics might both have need of a particular type




Figure 3. Schematic representation of how CPA allocates the processors in an MP machine among the jobs queued for processing.

of Trilinos solver package that could be “plugged-in” to each via a standard API. Hence, Trilinos offers the developer the opportunity to greatly reduce programming time. And as a corollary, Trilinos improves the efficiency with which Sandia’s MP computing resources are applied to crucial national security issues, such as stockpile stewardship, climate change, energy security, and others.

Other than the pre-existence of Aztec, a version of which ultimately became a package within Trilinos, subsequent developmental activity that spawned Trilinos continued in 1998, with an LDRD project that produced Epetra, what came to be the most important package in Trilinos, a cornerstone parallel linear algebra solver that provides the data structure for Trilinos but also can be utilized independently of other Trilinos packages. Its importance is underscored by Heroux’s own assessment: “It can be safely said that Trilinos would probably not exist if not for that LDRD.” Although the Advanced Simulation and Computing (ASC) Program has served as a key and significant funding source for the ongoing development of Trilinos, particularly its adaptation to new hardware platforms as MP computers become even faster and more powerful (sustained petaflop processing, with exaflop processing in the offing), ASC has “rarely had the luxury of high-risk funding.”

This observation succinctly defines one key reason for the importance of the LDRD program in this arena.

This research successfully yielded not only Epetra, but another solver known as GOMA, a multiphysics fluid dynamics solver that eventually evolved into ARIA, a molecular dynamics code that assists scientists and engineers in understanding the interactions among the particles in fluids, a key issue in nanoscience. A second Trilinos-associated LDRD broached the issue of data repartitioning in parallel solvers, combining a set of data repartitioning tools in a package known as Zoltan and implementing it within Trilinos. Zoltan rebalanced the workloads of the different processors in a parallel simulation, more or less “on the fly” as processor requirements changed during the course of a simulation running on MP hardware.

For example, Trilinos has been essential to the 2008 R&D100 award-winner Xyce™, a software package that is heavily relied upon to model electrical circuitry at Sandia—some of this circuitry key to critical weapons systems. At the same time, these software libraries have been collaboratively applied to industrial problems by companies such as Boeing and Goodyear Tire and Rubber.

OPTIMIZING HARDWARE ALLOCATION

In the context of the ongoing growth in hardware power, Compute Process Allocator (CPA), derived from a Sandia experimental MP computer known as CPlant (fig. 1), captured a 2006 R&D100 award for its ingenuity in hardware optimization (fig. 3). For as the hardware becomes more complex, the rate of individual processor failure will be unlikely to improve, and may, in fact, become even more of an issue. CPA uses algorithm theory to improve the throughput of jobs run on MP machines.

Based on Sandia's close affiliation with Cray Inc. in the implementation of Sandia's Red Storm supercomputer (fig. 4), a Cray XT3 hardware platform, Sandia staff—in collaboration with academic partners—applied CPlant lessons to the Cray XT3 hardware in the form of CPA. CPA utilized mathematical algorithms to optimize processor allocation to given jobs simultaneously running on the hardware. Specifically, it uses both a space-filling curve and span minimization to ensure that it optimizes the allocation of processors running a particular job. The developers demonstrated that this improved system throughput by 23%, that is, 23% more jobs could be simultaneously run during a given time period. This ultimately represents not only a facilitation of research efforts requiring MP computer time, but also a cost-savings, as well, given the high operating cost of such hardware from both electrical energy and staff-hour perspectives. CPA has impacted work in global climate change, nanoscience, astrophysics, and military applications.

MEMORY BOTTLENECKS

This incredible processing power brings to the forefront a problem best termed the memory bottleneck. In most cases, computations on data can be executed as fast as the data can be moved from remote memory to local memory, and from local memory to their computational processors. A failure to move data with sufficient rapidity means that the processors, super-fast and numerous though they may be, may be starved for data upon which to compute. It is as though we have increased the horsepower by adding cylinders to an auto engine, but failed to upgrade the fuel injection system that moves adequate fuel into the larger and more-powerful engine. Intel, Inc. (Red Storm's processor manufacturer) is employing LDRD research that is developing network



Figure 4. Cabinets containing racks of processors for Sandia's Red Storm supercomputer.

methodologies to deliver data to processors, which would otherwise be “data starved.”

All told, the impact of over 15 years of LDRD-supported research into MP computing has progressively moved forward Sandia capability in this arena, which has, in turn, improved the speed, accuracy, and reliability with which a variety of complex challenges can be approached. From climate change to critical nanosystems and electrical circuits to stockpile stewardship, this work has illustrated a quite palpable payoff from LDRD investments. ■

Designing Clever Molecular-capture Cages for Nonproliferation and Sensing Critical Environments

Using Metal Organic Frameworks (MOFs) can allow spectral-based detection of SNM for nonproliferation—and of gases within weapons




Figure 1. A simple representative MOF.

Picture yourself shrunk down to the nanoworld; no, not the microworld of submarining through blood vessels and doing surgery in the brain, one nerve cell at a time, but smaller, much smaller. Water molecules are familiar, most protein molecules gigantic. Your task: design nanocages to capture, collect—and identify—particular substances, or even particular types of radiation.

MOFS: NANOLABORATORIES FOR MOLECULAR ENGINEERING

Beginning with LDRD funding in 2006, Sandia chemists and engineers have been studying metal organic frameworks (MOFs) because of their controllable nanoporous properties. What this means is that a bonded combination of a metal ion together with organic molecules can form a diversity of molecular structures with controllable nanopores—nanocages, usually of very low density because the pores (cavities) formed by the 3D molecular geometry create more space than substance. By varying the specific chemical constituents, the molecular geometry of the cavities can be controlled, leading to nanoscale cages of different sizes and shapes that can trap or sequester a variety of other substances within those nanocavities, in some cases, the ancillary substances “accessorizing” the MOF toward particular functionalities. MOFs offer unique opportunities to control pore size, surface area, and chemical properties not readily accessible in other material classes. This gives MOFs numerous potential applications in gas sorption (including the problem of carbon dioxide sequestration), electro-optical-

chemical sensors; separation methods; and permeable membranes for fuel cells. This article considers two MOF applications pertinent to the weapons program. The first, in the area of nonproliferation, enables a new type of sensitive detection of special nuclear materials (SNM), the second, as sensor-monitors of the health of weapons themselves.

A simple exemplary MOF is shown in figure 1 — Zn_4O

(the metal ion, shown as blue polyhedra) joined by benzene dicarboxylate linkers (the organic portion; oxygen, red and carbon, black) to give an extended 3D cubic framework with interconnected pores of 1.1-nm aperture width and 1.8-nm pore diameter (the yellow sphere represents the pore). To get a feeling for the tiny scale of these molecular cages, it may help to recognize that the size of an average bacterial cell is about 1 μm (or 1000 nm), immensely larger than these nanopores.

A slightly different MOF (named MOF-S1, also based on zinc but with longer linkers) helps to illustrate the cage-like structure in both extended and space-filling renderings. In this work, the organic linkers comprising the nano cage (fig. 2) emit light when exposed to radiation (see below).

USING COLOR TO DETECT IONIZING RADIATION

Detecting radiation emitted by fissionable material and other radiation sources is of key concern to agencies involved in homeland security and treaty verification. One way of doing this is to use molecules known as scintillators, which emit short pulses of light when high-energy particles collide with them. The major challenge is to distinguish high-energy neutrons—a signature for weapons-grade plutonium—and thousands of gamma particles that are also emitted by these




Figure 2. MOF S-1, a stilbene-based MOF retaining scintillation properties of the parent molecule, stilbene.

sources, as well as common commercial radiation sources such as smoke detectors and medical isotopes. This can be done by recording the intensity of the scintillation light as a function of time, which has a shape that depends on the particle type. Such measurements are not simple, requiring sophisticated electronic equipment to monitor the light output of the scintillator from very short times (nanoseconds) to much longer times (microseconds). Another problem is that the best scintillators are liquids that are both flammable and toxic.

Beyond the fact that such liquid scintillators are hazardous, comes an issue with their sensitivity. Because of rules that govern how electrons can pair up with each other, only about 25% of the possible light is actually emitted. Unfortunately, signal is often at a premium in radiation detection, because the source is weak, heavily shielded, or at a long distance from the detector. It is thus critical to produce as much light as possible for every particle that the scintillator interacts with.

Enter a MOF with an organic component derived from a well-known scintillator called stilbene (fig. 2). This MOF locks the stilbene into a particular molecular geometry that causes it to emit more light. The idea for using this material for radiation detection was stimulated by discoveries in Mark Allendorf's




Figure 3. Bi-spectral (blue & red) luminescence by crystals of a DUT6 host MOF infiltrated with the Ir(quin)₃ guest (which fits into the pore indicated by the purple sphere).

lab during an LDRD-supported MOF project. The fact that the stilbene in the MOF retains its properties as a scintillator led to the realization that a large number of MOFs can behave as scintillators, creating the first new class of such materials in 50 years. Moreover, the MOFs' crystalline structure is a kind of "nanolaboratory" that creates numerous opportunities for clever molecular engineering of a broad range of photophysical processes.

A particularly exciting realization was that MOF pores could serve as hosts for "guest" molecules that could enhance the scintillation process. Taking a cue from the world of solid-state lighting, the Sandia researchers filled the MOF pores with small molecules containing heavy metals such as iridium. These compounds, known as triplet harvesters, are used to make bright organic light-emitting diodes used in cell phones and other electronic displays. When applied to radiation detection, they efficiently "harvest" the 75% of the excited electrons that in a conventional scintillator produce very little light. This means that more light is emitted, which can speed up detection and reduce the rate of false positives.

Even more intriguing, the triplet harvesters can be used in a new way to detect radiation. Instead of measuring the radiation-induced light as a function of time, the color of the light can be used. The triplet harvester and the MOF emit light

at different places in the visible-near ultraviolet part of the spectrum. The MOF typically emits blue light, while the emission of the triplet harvester can be tuned from the green to the red. For example a MOF host known as DUT-6 (fig. 3) combined with the guest Ir(quin)₃ will emit spectrally distinct signals at around 380 nm (blue-violet) and 600 nm (red), and the characteristics of these signals as elicited by absorption of a neutron as compared with absorption of a gamma ray can form the basis for spectral shape discrimination (SSD) of neutrons and gamma rays. The host MOF and guest molecule can also be selected to emit light at matching wavelengths for use as higher-sensitivity detectors in the more-traditional time-based detection approach described above and known as pulse-shape discrimination (PSD) (fig. 4). Different particles lose energy at different rates when absorbed by a material: gamma rays lose much less than do neutrons. This is the reason that PSD works, and for SSD—using the MOF, DUT-6 and the guest, Ir(quin)₃—one gets more blue emission from a gamma ray compared with a neutron. The ratio of the blue light intensity to the red light intensity defines the incoming ionizing energy as neutron or gamma. This method is far easier to interpret than the often complex nuances of PSD.

While attending to the constraint of MOF host-nanocage size and guest molecular dimensions, one can see why MOFs provide the fodder for nanolaboratory molecular engineering to create substances for numerous specific applications. In effect, by funding basic research on MOFs, the LDRD Program opened the door to a vast range of mission-pertinent applications for this chemistry. "This is an example of how LDRD should work," comments Allendorf.

Indeed, subsequent to this round of LDRD funding came funding from the Defense Threat Reduction Agency (DTRA) and DOE to develop MOF-based scintillators. Accurate detection and characterization of ionizing radiation forms the basis for the ability to identify illicit transport and manipulation of SNM, and therefore, is critical for US nonproliferation initiatives.

SENSORS FOR SURVEILLANCE OF THE ATMOSPHERE WITHIN WEAPONS

A second application being actively pursued in collaboration with Georgia Institute of Technology (Georgia Tech) is the

use of MOFs to sense the gas atmosphere within a nuclear weapon in real time, changes that can indicate alterations in components (an altered “state of health”), but that can also be used to validate models that simulate the weapon environment. The physical basis for such gas detection is the static microcantilever (fig.5), a structure about one-fourth of a millimeter (250 μm) long that can be implanted in the interior of a weapon in the form of an electronic module about the size of a USB thumb drive.

The principle here is that a minuscule change in the position/shape of the microcantilever can be electronically sensed as indicating some selected change in the environment of the device. In other words, were one to place 1000 molecules of substance X on the microcantilever, its structure might shift sufficiently to provide a signal (imagine the bending of a microscopic “diving board”), thus acting as a sensor for substance X when fitted with electrodes able to sense the physical change in the microcantilever. Static microcantilevers can have sensitivities of picograms (10^{-12} or one-trillionth of a gram of bound material). The question then becomes, how does one “functionalize” the surface of the microcantilever, so that it will selectively bind substance X, and not other molecules that might be present. The solution is to coat the microcantilever’s surface with a MOF chemically tailored to absorb the gas that one desires to monitor. With the amazing range of structural motifs and nanocage sizes available in the MOF family, this appears to be eminently feasible, particularly with simulations to assist the experimental work.

SUCCESSFUL DETECTION

This work has already demonstrated the successful detection of humidity—water vapor—at sub-parts-per-million (ppm) levels by means of microcantilever sensor operating over a ten-month period. Many other gases can now be detected as well. Changes in some of these gas concentrations are important because outgassing from weapon components can indicate aging-related and other changes in weapons’ state of health.

With all the already remarkable applications discovered for MOFs, it is nonetheless likely that there is much more to come, and that many Sandia missions will be beneficiaries. ■




Figure 4. Schemes for detectors employing photomultiplier tubes (PMT) and filters to conduct either PSD- or SSD-based detection of special nuclear materials.




Figure 5. Electron micrograph of microcantilever, prior to coating with a MOF.

“This is an example of how LDRD should work.”

Radiation Hardened Electronics for Critical Systems

Protecting the functionality of integrated control circuitry in high-energy environments




Figure 1. Sandia rad hard ASIC chip.

“If such control systems operate in high-energy environments . . . the only way to ensure functionality is by rad-hardening.”

Since the advent of atomic weapons in the 1940s and of orbiting satellites a decade later, electronics designers have had to face the issue of how to sustain operation in the presence of damaging radiation to delicate, progressively miniaturized electrical components. The nomenclature for this operational resistance to high-energy (or “ionizing”) radiation—such as the gamma radiation in space—is “radiation hard” (abbreviated “rad-hard”), an indication that electronics are radiation-damage resistant and can continue to operate properly even in a high-intensity radiation environment.

DAMAGING OPERATIONAL ENVIRONMENTS

Application-specific integrated circuits (ASICs) are electrical components designed to function in specific electronic systems, and they are ubiquitous in both military and civilian control systems (fig. 1). If such control systems operate in high-energy environments, or if they may at some point be subject to such conditions, the only way to ensure functionality is by rad-hardening. For example, satellites in orbit inevitably accumulate a large dose of gamma radiation, neutrons, and charged particles from the sun over their operating lifetime. Nuclear weapons both generate such high-energy conditions and are potentially subjected to them. A nuclear explosion can send a pulse of radiation (in the form of gamma rays




Figure 2. Colorized images of the Jovian moon, Io, taken by the spacecraft, Galileo

and neutrons) large distances from the explosion site, thus affecting the function of any ASICs in proximal control systems. In addition, during their operation, such weapons would very likely be subjected to defensive actions, including high-radiation measures by adversaries attempting to disable their control systems, which must therefore be rad-hardened. Military sensors in the field can likewise be subjected to radiation during adversarial attempts to corrupt their functionality.

The common feature of these high-energy environments is the ionization of atoms and molecules in the component material of the ASICs, both directly changing electronic structure and creating holes (or sinks) that attract electrons from other molecular components. Wherever it impacts the atoms of an absorbing material within an ASIC, a high-energy neutron, for example, can wreak havoc, successively ionizing (knocking electrons from) atoms, leaving a trail or collection of electrons that can flow through the material to switch on a transistor, open an electronic gate, change a resistor's value, add an unexpected "bit" to a computer calculation, ablate a

memory, or short out a capacitor. In general, passive electronic components can change their electronic values or short-out via current flows initiated by the electrons freed through such ionizations, while active components such as transistors can switch from off to on or vice-versa. ASICs composed of millions of transistors can easily cease functioning entirely or are highly likely to miscommunicate critical control or sensor information.

SANDIA'S LONG-TERM COMMITMENT TO RAD-HARDENING

Since the mid-1960s, Sandia researchers have worked to reduce the sensitivity of control systems, communications systems, firing sets, and sensors to such damaging radiation impacts, thus preserving control integrity under these extreme conditions. While LDRD projects have only impacted rad-hardening since the LDRD program's inception in 1992, radiation hardening of electronics was actively pursued at Sandia, largely in secret, during the 1960s and 1970s. The DoD/



Figure 3. Artist's conception of the Mars Rover.

US military was the Laboratories' primary customer and worked closely with Sandia and several commercial vendors to find new materials and to understand radiation-damage physics. Originally, the simplest approach to rad-hardening was to pre-irradiate electronics to attempt to saturate the damage level, that is to pre-damage the material while retaining functionality. Such treatment reduced circuit gains, but it also was reasonably effective at limiting further damage. Since most commercial electronics manufacturers were selling to the military, they worked closely with Sandia to rad-harden their products. However, by the mid-1970s, as very large scale integrated circuitry was entering the landscape, the consumer electronics market exploded. Consequently, by the early 1980s, commercial vendors had turned away from the military market, which was now only a minor portion of the economic landscape.

Today, only Sandia and Honeywell remain contributors to the radiation hardened electronics manufacturing market, developing new solutions to reduce component size to the nanoscale and increase capability in high-dose and pulsed-radiation environments. Sandia supplies rad-hardened components to customers, such as NNSA, NASA, the Navy, and the Air Force for both aerospace and weapons applications. Sandia purchases commercial electronics engineered to customer specifications or, using electrical-circuit simulations, designs and fabricates rad-hard ASICs for its various customers.

MATERIAL SCIENCE INNOVATIONS AND APPLICATIONS

To thwart radiation damage, Sandia LDRDs have employed fundamental materials science initiatives to study different materials such as silicon-on-oxide, gallium arsenide, germanium, and chalcogenides, materials that restrict the flow of electrons freed by gamma radiation and neutrons, and that therefore limit the damage when compared to silicon—the standard material for microelectronic circuitry. Some materials, like the chalcogenides (sulfur, selenium, and tellurium exemplify chalcogen elements) use physical phase changes rather than charge storage for memory retention. Other solutions include the use of optical rather than electronic interconnects to maintain intra-circuit communications.

Along these lines, a significant Sandia contribution to the space program is exemplified by the NASA spacecraft, Galileo, which carried out an investigation of the planet Jupiter and its moons during the 1990s (fig. 2). Sandia designed, fabricated, and tested the radiation-hardened integrated circuits on the spacecraft, which returned enormous quantities of data from highly challenging radiation-filled environments such as that of the Jovian moon, Io, to which Galileo passed in rather close proximity. The Mars Rover (fig. 3) exemplifies another spacecraft requiring rad-hard ASICs to enable functionality in the high-radiation environment of space.

The launch of the entirely Sandia engineered Multispectral Thermal Imager (fig. 4) in 2000 is illustrative of the Laboratories' and LDRD's contribution to US and international space initiatives with rad-hardened electronics. A space package that was commissioned to fly for three years, primarily as a radiometric and thermal detector of nuclear proliferation, it has flown for about a dozen years. In addition to monitoring the atmosphere for evidence of activities involving nuclear materials in the interests of nonproliferation, it has returned data regarding valuable geophysical and environmental monitoring and assessment, in areas such as volcanoes, glaciers, climate change, and studies of the moon.

C-RAM AND RAD-HARDENED INSULATORS

LDRD projects have continued to contribute new concepts for rad-hardened computer memories, sensors, interconnects,

and capacitors, primarily, although not exclusively applied to weapons and avionics. A 1996 project addressed processing techniques for silicon-on-insulator memories that were based on the presence of defects to reduce the transport of radiation-induced charges through the material. Such memories have subsequently become commercial off-the-shelf products. During 2003-2004, a new concept was examined to produce highly radiation-hardened memories using chalcogenides, for example, mixtures of germanium, antimony, and tellurium. These devices based memory upon phase changes in crystal structure that immensely increased local electrical resistance, to create a computational “bit” of stored memory. The polycrystalline state (low resistance) provided the “zero” and the amorphous state (~1000-times greater resistance) provided the “one.” These devices were coprocessed with the Air Force Research Laboratory, and today this C-RAM (chalcogenide random access memory) is another off-the-shelf product, an outcome of Sandia LDRD-funded research.

Another LDRD project addressed the radiation sensitivity of capacitors, a ubiquitous charge-storage component in ASICs. Such charge stored in capacitors can be easily and inappropriately released (discharged) by gamma rays, neutrons or other energetic charged particles, so the project developed a new insulator that would permit high stand-off voltages while having extremely limited electrical current flow when exposed to radiation sources. Beginning with Mylar, a commonly used insulator, but one that was not radiation resistant at high voltages, project staff doped Mylar with trinitro-fluorenone (TNF)—which had been previously used to produce coatings on auto windows for sunroofs—thereby dramatically increasing its rad-hardness. And when TNF became unavailable after its reclassification as an explosive, the team had to resort to clever chemical engineering. By devising novel synthetic techniques, they still managed to dope Mylar and demonstrate significantly improved rad-hard capacitors fabricated in-house.

One major outcome of the LDRD investment in the science and technology of radiation hardening is the global recognition of its capabilities leadership. For example, in the late 1990s, the Intel Corporation chose Sandia to rad-harden its Pentium computer processor for space and defense applications, in simultaneous recognition of both national need and Sandia’s unique capabilities. Recently, Sandia has collaborated with IBM, Boeing, and Honeywell in a DTRA program investigating rad-hard computing technologies for space applications.



Figure 4. Artist’s conception of the Multispectral Thermal Imager in orbit above earth

Rad-hard LDRD projects have provided creative opportunities for LDRD researchers to directly contribute to the nuclear weapons, space programs, and defense systems arenas, as well as to fundamental chemistry and chemical engineering research. These investigations seeking continual improvements in the areas of rad-hardening of integrated circuits represent advances in materials science that will undoubtedly continue to manifest positive outcomes in a variety of national security areas. ■

Ingenious Lubrication Schemes for Preventing Wear in Critical Electromechanical Components

Paving the way for improved lubrication and further miniaturization of critical components




Figure 1. The formulaic and space-filling models of the amphiphilic 1-octanol molecule showing its hydrophobic (nonpolar) and hydrophilic (polar) portions.

“... the design space is once again open as far as your imagination.”

Lubrication is usually necessary to keep mechanical components functioning smoothly and to prevent excessive wear at points of contact, where friction and adhesion forces can act to adversely affect performance. As weapon safety and security devices are scaled down to save weight, reduce volume, or provide added functionality, existing approaches for lubricating their parts become impractical, and may even damage parts. A new lubrication approach circumvents these problems.

PURSuing NEW LUBRICATION METHODS

Currently, Sandia critical electromechanical systems for both use control and environmental sensing undergo a careful process of surface preparation, masking, and lubrication with solid lubricant coatings using traditional methods such as painting on with an organic binder, or impingement of solid particles on the surface. Although adequate for these mostly stainless steel surfaces, there are reasons to rethink this methodology as these systems evolve. First, these systems function in controlled environments where size and weight are significant considerations, and therefore, the imperative to downsize is always a factor. And as discussed in more detail below, for microelectromechanical systems (MEMS), the parts are too small for these traditional solid lubricants, and the chemisorbed films used to reduce surface energy wear

out quickly during contact. Second, to lubricate in these traditional ways requires preparation of stainless steel surfaces such as brushing and burnishing, and as parts get smaller, they also tend toward increased fragility. It is therefore important to minimize handling steps that risk damage. Hence, what is desirable from both these standpoints—tight tolerances and minimization of manipulation—is the ability to lubricate using a technique that would apply a very thin yet effective layer to surfaces through as gentle a procedure as possible.

THE ALCOHOL SOLUTION

Motivated by these necessities, Sandia's Mike Dugger proposed an LDRD-funded project to examine these phenomena more closely in an attempt to develop a chemistry that could lead to the lubrication of surfaces from the gas phase during operation, with minimal manipulation of the electromechanical component requiring such lubrication. Beginning in 2006, a Pennsylvania State University doctoral student was studying adhesion of silicon surfaces using atomic force microscopy (AFM) to measure the force required to pull the instrument's tip from the surface and discovered that alcohols reduced that requisite force. Upon joining Sandia as a student intern, the student and Dugger began studying the surface reactions when silicon surfaces were rubbed in the presence of alcohol vapors. They discovered that at the points of contact, a whole host of complex cyclic organic products were created. They hypothesized that the heat generated by frictional contact, possibly resulting in thermionic emission of low energy electrons from the surface, broke bonds in the adsorbed simple organic molecules and allowed more complex species to form.

Given their amphiphilic nature, alcohols were a reasonable starting point. Amphiphilic molecules have both a hydrophilic (or polar) portion and a hydrophobic portion and linear alcohols fit that bill, their polar alcohol (OH) group with hydrophilic properties and their nonpolar hydrocarbon chains of covalently bonded carbons and hydrogens hydrophobic. For example, 1-octanol, $\text{CH}_3(\text{CH}_2)_7\text{OH}$, obviously has that amphiphilic property (fig. 1), and would be expected to bind to metal surfaces by its hydrophilic end, leaving its hydrophobic end exposed as a lubricant. The series of linear alcohols — from the 1-carbon (C) methanol through the 2-C ethanol, 5-C, pentanol, and 8-C, octanol — provided lubricant possibilities. But in addition to having the required




Figure 2. Wear debris on a gear operated in an N_2 environment (top) compared with no detectable wear debris on the same gear type operated in pentanol vapor.

amphiphilic/lubricant properties, Dugger recognized that there were other considerations: alcohols are flammable, such that safety was one issue; additionally, since prior research at Sandia on vapor-phase lubrication (VPL) of silicon surfaces suggested that the least-intrusive and gentlest way to introduce an alcohol into a system would be as a vapor (gas), he would need to work with one that had a convenient vapor pressure regime over the temperature ranges at which critical electromechanical systems would be expected to function.

In the end, the alcohol with the required properties was the 5-C, pentanol, and while testing related compounds, such as pentane and thiols—organosulfur molecules, where the alcohol (-OH) group is replaced by the thiol (-SH) group—the research, in collaboration with Professor Seong Kim of Pennsylvania State University, set out to optimize chemistry for lubricating both silicon and stainless steel surfaces. The project also tested various pentanol-water mixtures to ascertain how much water could be tolerated and still provide lubrication (water vapor competes with pentanol vapor for adsorption to surfaces requiring lubrication). A key parameter to determine was the pentanol vapor concentration that would result in a monolayer of pentanol molecules on those surfaces requiring lubrication. The result of such pentanol delivery was quite convincing in terms of successful lubrication. Components allowed to run many cycles without such lubrication showed marked wear, debris accumulation, and deterioration, while the same systems run with pentanol




Figure 3. Sub-millimeter-sized intermeshing gears in a MEMS device.



Figure 4. Comparison of a MEMS switch (top) run in air (middle)—showing wear-product buildup—and in a pentanol vapor environment (bottom), showing almost no buildup.

lubrication were essentially debris and wear free (fig. 2). A key finding was that none of the components experienced adverse alterations in their electrical properties.

MAINTAINING FUNCTIONALITY IN MEMS

In addition to providing an alternative—and greatly advantageous—lubrication methodology for extant systems, essentially a new paradigm, VPL has introduced the possibility of even greater size and weight reduction, by its ability to effectively lubricate microelectromechanical systems (MEMS) with moving, interacting parts on the order of fractions of a millimeter in size (fig. 3). Frequently made of metals such as stainless steel or nonmetals like silicon, such microscopic gears, switches, and gates are usually hydrophilic, attracting water vapor that initiates oxidizing deterioration chemistry, and although attempts have been made to coat these surfaces with hydrophobic coatings, evidence shows that these coatings are relatively easily damaged by contact between and among surfaces—such as intermeshing of gears or sliding of two surfaces along each other. Added to this is the fact that attempting to apply traditional coatings to such tiny components even more emphatically emphasizes the change in tolerances because even a thin layer would still be thick enough to change dimensions when measured against such minuscule parts. Because of the size and weight reductions that are desirable in many mission-critical systems, miniaturization of critical systems to include MEMS is a desirable goal for next-generation systems. This LDRD-funded research speaks to that future.

Adhesive forces between contacting parts—such as electrostatic, van der Waals, and

capillary forces—become much more significant due to the small scale, and they tend to overwhelm the small actuation and mechanical restoring forces for MEMS devices. The moving parts therefore freeze and fail to move. (By analogy, imagine trying to slide two smooth surfaces placed in contact with a few drops of water between them. The water capillaries would “glue” them together, inhibiting movement.) Thus, in order to have reliable, freely operating MEMS, the contacting surfaces must be lubricated in a manner that minimizes these forces. Consequently, what is required as such electromechanical systems continue to shrink in size is a method for applying a molecular monolayer of some substance, such that electrostatic and other adhesion forces are reduced or canceled—the smaller/thinner the introduced molecular layer, the better.

The initial enthusiasm for MEMS applications in the 1960s and 1970s was gradually eroded by the realization of the difficulty posed by such forces. Hence, for the past 30 years, MEMS have mostly proven useful for applications with no contact or sliding, given that these adhesion forces always present problems. High-end applications include pressure sensing (catheters), bio-MEMS (chemical and biochemical analysis and point-of-use diagnostics), infrared imaging and microphones. MEMS also appear widely in consumer products: inkjet printers, the accelerometers that deploy your car’s airbags, and the tiny gyroscopes that sense the movement of the way you rotate your iPhone or iPad. But MEMS could also be quite important to the functioning of surety systems and electromechanical locks for information protection/access, particularly if a new generation of smaller size and weight components becomes reality. The ability to create this functionality with many moving parts at a

size scale of millimeters or below means that safety and use-control systems can be implemented with much lower mass and volume.

It is worth noting that a liquid lubricant wouldn't work for MEMS because the viscosity of the liquid is so high, compared to the very small force that actuates a MEMS device. (Moving the MEMS device in the presence of liquid lubricant has been compared to trying to swim through a swimming pool full of honey.)

In the end, pentanol vapor lubrication proves to be a solution to this problem as well. A key parameter to determine was the pentanol vapor concentration that would result in a monolayer of pentanol molecules on those surfaces requiring lubrication. The result of such delivery was quite convincing. MEMS systems allowed to run many cycles without such lubrication showed marked wear, debris accumulation, and deterioration, while the same systems run with pentanol lubrication were essentially debris and wear free (fig. 4).

DELIVERY SYSTEMS

A key additional goal was to develop delivery systems for introducing pentanol vapor, and to assess the performance and durability of the lubricant. If a system could be devised that released pentanol from a polymer contingent upon temperature changes, this polymer “outgassing” could automatically deliver requisite amounts for ongoing lubrication. Precursor polymers could be synthesized such that, upon exposure to high temperatures, small molecules like pentanol would be released. The on-demand delivery of vapor-phase pentanol would allow for scheduled or as-needed lubrication of critical components—including, ultimately, MEMs, greatly improving their lifespan. Central to the design of this on-demand system was the strong (covalent) attachment of the lubricating molecules to the polymer backbone, ensuring that the lubricant would only be released at high temperatures ($> 145\text{ }^{\circ}\text{C}$). Because most systems would not see ambient temperatures near the release point of the lubricant, the lubricant pentanol would remain attached to the polymer until needed. At that juncture, pulsing the temperature to $150\text{ }^{\circ}\text{C}$ would release a pulse of pentanol lubricant into the system.

Toward this end, project staff synthesized a pentyl-xanthate polymer for such “on-demand” release of pentanol. Since that polymer also released carbon disulfide when heated, an alternative polymer was synthesized that released only pentanol. This resulted in a successful system for such polymer-outgassing on-demand delivery for most applications of MEMS technology, which would function in response to precise pulse heating to evolve pentanol vapors. Unfortunately, certain military applications were deemed more difficult for this methodology because of the rapid temperature shifts that they experienced during operation.

Several epoxy materials were also prepared that contained pentanol for vapor delivery at elevated temperature. The project team used thermogravimetric analysis to demonstrate that pure pentanol vapors could be quickly generated at elevated temperature. The research concluded that pentanol vapors could be successfully generated from these materials (fig. 5). Alternatively, during manufacturing of components in a controlled packaging environment, the devices can be packaged in the presence of the requisite vapor pressure of pentanol, thereby effectively pre-lubricating them.

Because of the simplicity and effectiveness of VPL, both current and future paths now appear far more open. For example, according to Dugger, “people have shied away from MEMS solutions because of those problems, but the design space is once again open as far as your imagination.” He hopes that he will see the day when that design space is applied to MEMS solutions within Sandia's national-security mission space. ■




Figure 5. Synthesis of a polymer that will outgas pentanol vapor at $150\text{ }^{\circ}\text{C}$

Encapsulating and Protecting Delicate Electronics

The chemical engineering and modeling underpinning polymeric foams encompass a complex interplay of challenges




Figure 1. Micrograph of an epoxy foam illustrating its structure as a collection of bubbles trapped within a polymeric framework.

“challenging materials where everything has to be perfect.”

Sometimes, in researching a topic that appears from a distance to be reasonably simple and straightforward, one is ultimately compelled to reconsider that judgment as naïve. The topic of foams turns out to be just such an instance.

Whether it’s the packaging foam in your UPS shipment or the protective enclosure for a delicate bit of electronics within a weapon or other critical system—a so-called chemically blown (expanded) foam is produced by a liquid mixture of organic molecules poised to undergo a set of chemical reactions that will first, generate a gas (often CO_2) that creates bubbles and expands the liquid into a foam, then stabilize and harden (or “cure”) that foam into a solid by inducing the forming of chemical bonds that complete the polymer of which the solid foam is ultimately composed. Whether epoxy-based or polyurethane-based—the two predominant polymeric chemistries—the goal is a solid of a specific shape with a small and uniform bubble size (fig. 1) that resists compression, able to tolerate external pressure and other forces while protecting some structure that it encloses or surrounds.

In addition, for structural foams, the dimensional stability, or lack thereof, with aging must be understood. Foams can also be physically blown using a volatile blowing agent, in which case, the foaming reaction is initiated/controlled by altering some physical parameter, such as temperature or pressure.

A LONG HISTORY OF PROTECTING ELECTRONICS

Foams play a role in several areas of Sandia’s mission space, but none quite as important as protecting delicate electronics

from physical insults, thereby guaranteeing their ongoing functionality. Although they have successfully done so for decades, legacy foams suffer from certain drawbacks, including the inclusion of chemistry that can immunologically sensitize workers to re-exposure (a form of allergic reaction), as well as being difficult to remove once in place. Additionally, physically blown foams employed in certain Sandia applications, are more difficult to control.

Given these circumstances, the Sandia LDRD program funded research into a study of alternative foam approaches from both a chemical engineering and a modeling perspective. If anything, this project validated the view of Sandia chemist Mat Celina that foams are “challenging materials where everything has to be perfect.”

TIMING IS EVERYTHING

A foam takes its shape by being poured or forced under pressure into a mold, where, as a reacting liquid that expands as gas bubbles forms, it must fill all spaces—sometimes extremely narrow ones. Then, just at the right moment, the foaming process is stopped, and the foam cures by polymerizing the liquid into a solid. This creates a homogeneous foam with evenly spaced bubbles of a consistent size, and therefore consistent protective properties as it forms around a delicate bit of electronics as defined by the mold’s size and shape. (This “right moment” description is idealized, given that, in real situations, the foaming and curing reactions are occurring during the same time window). If this curing occurs too slowly and/or the foam-liquid’s density is too high, gravity can drive un-bubbled liquid to the bottom of the mold (fig. 2A). In other situations, such as in delayed curing of foams of lower density, bubbles can merge together creating an unevenness in bubble size and spacing (fig. 2B). This phenomenon known as coalescence, where bubbles merge together, and Ostwald ripening by which gas moves from smaller to larger bubbles, also can contribute to a lack of homogeneity in bubble size and distribution. Such undesirable attributes can leave a foam with weak spots, and thus, with inadequate protective properties, inadequate resistance to compression and other stresses. Conversely, if the curing (polymerization to a solid) occurs too rapidly, prematurely eliminating liquid properties (rheology) of the foaming chemicals, the foam will fail to fill all spaces in the




Figure 2. A: Drainage of unfoamed liquid at the bottom of a mold because of high density and a delay in curing. B: Overly large bubbles formed during curing of a lower-density foam.

mold—particularly very small, narrow spaces (see fig. 3), leaving voids without foam, a clearly undesirable outcome.


PI Lisa Mondy and her LDRD research team studied a diversity of new routes to such chemical foaming with production of CO₂, chemistries that would reduce toxicity and bring a greater measure of control over and understanding about the foaming and curing processes. With a background in multiphase flow, Lisa began her Sandia career with a study of geothermal drilling foams, and after several successful proposals to study the flow of suspensions, she began working with Rekha Rao in modeling particle migration, and this collaboration brought the pair to encapsulation foams through some work with Sandia staff encapsulating critical components with particle-filled epoxies. Assisted by campaign and Advanced Simulation and Computing funding, they began to model extant stockpile foams. In the process, the problems they encountered provoked creative ideation that led to funding of the LDRD project and collaboration with creative material chemists like Celina and manager Jim Aubert, a chemical engineer who worked on foams early in his career. Celina, whose background lay in applied polymer science and materials optimization had previously worked with the Kansas City Plant on reliability issues in foams processing.



Figure 3. Comparison of foams within mold, relatively homogeneous foam, at left, compared to a foam with voids (at right) a result of bubble coarsening. Filling of narrow passageways such as the top horizontal beam of the mold is always difficult, and is somewhat easier if foam curing can be delayed slightly to maintain foam liquidity.

As part of the LDRD, the team's modeling of the existing chemistry (fig. 4) was guided by the long-term goal of extending this model to whatever chemistry appeared to be the most propitious. Ultimately, the intent was to render the model adaptable to diverse chemistries, a breakthrough engineering model that did not previously exist. With respect to the experimental chemistry, the goals of the project were to move toward the best aspects of existing foams without their liabilities, and to project potential future needs of life extension programs. This aspect of the project involved a collaboration with the Kansas City Plant, a collaboration that is ongoing.


After studying numerous methodologies for generating CO₂ bubbles, the project developed novel, stable epoxy foams using a liquid epoxy anhydride system that generates CO₂ bubbles through the decomposition of tert-butoxycarbonyl anhydride, the reaction generally representable as follows:



CHALLENGING CHEMISTRY

One goal was to develop a system in which the foaming and curing reactions were more separable than the equivalent reaction in extant polyurethane-based foams, in which the foaming and curing reactions are intrinsically coupled to a greater extent. Such coupling makes it more difficult to maintain the foaming chemistry in a liquid state in order to fill difficult narrow passageways within a mold, the existing foam chemistry characterized as “unforgiving.” “Foams are tough systems; all the chemistry tends to be happening together,” says Celina. This was especially problematic in attempts to use spectroscopic signatures to follow the kinetics of individual reactions. Because of cross-talk in reaction chemistries, isolating a clear infrared (IR) spectroscopy signal for the rate of CO₂ evolution proved to be virtually impossible, particularly since additional CO₂ was evolved during curing, subsequent to that produced in foaming. Hence, despite a system, that did achieve a greater separation between foaming and curing, this was not complete, mostly because the cure process involves very effective chemistry that cannot be completely delayed.

Figure 4. Computational modeling of foam injection in a mold. The foam is modeled as a shear-thinning material, ignoring irreversible effects of bubble breakage. In large sections, the viscosity behaves as a Newtonian fluid and in small sections, the viscosity thins to a quarter of that value.



Ongoing support for the polymer chemistry and processing aspects of the research, for qualification in future weapons systems, would appear to be a judicious investment. Given the uncertainty regarding future needs and the consequent desirability for flexibility in being able to generate novel foams for critical life extension initiatives, such research looms as an important component for that future. Clearly, qualifying a new material in this particular arena is a rather difficult proposition. Part of the reason is the nature of the materials themselves. Starting from the non-intuitive proposition that generating a gas within a liquid ends up with the formation of a solid, foams—even with a high degree of knowledge and understanding about the individual chemical reactions involved—manage to behave somewhat differently than anticipated. A system that, at first glance, seems simple and straightforward, upon deeper examination, proves to be fantastically challenging. These initiatives that foster greater scientific and technological understanding should ultimately contribute to key national security requirements important to several Sandia missions. ■

More-reliable Methods of Verification for FPGA-based Digital Systems

Ensuring that critical control systems will perform as expected




Figure 1. The process of using vendor tools on design code to pattern the logic elements of FPGAs for critical systems may introduce unanticipated vulnerabilities that simulation would not necessarily account for.

n	Simulation run time
1	4 μ s
2	16 μ s
4	254 μ s
8	65 ms
16	1.2 hr
32	584,942 yr

Figure 2. Comparative run times for “complete” simulations of an n-bit adder.

*“how confident can I be that
what I want is what I get?”*

A fundamental question for Sandia high-consequence electronics involves the reliability of field programmable gate array (FPGA)-based control systems. Many critical-system electronic components are being replaced by modern digital devices, thus dramatically increasing system complexity. Such systems include those used in aerospace applications, as well as those upgraded as part of nuclear weapons life-extension programs, and the updated designs commonly rely on FPGAs to implement sophisticated logic. A key outcome of this trend is the need to validate the performance of such digital systems to the greatest extent possible.

FPGAs

Ubiquitous in modern hardware, FPGAs are electronic components that are somewhat akin to a tabula rasa, a blank page. Perhaps a more appropriate analogy would be a child’s play tool for improving reading—a magnetic slate filled with movable words that the child can reconfigure into a variety of meaningful sentences. In FPGAs, the words are logic blocks that can be connected together in a variety of configurations determined by a user who is addressing specific hardware design requirements. Based upon tools provided by the vendor, FPGAs are thus designed to be configured by a customer after manufacturing, that is, after

their logic functional blocks are implemented. One advantage of FPGAs over custom hardware is that they are harder to attack because the design for a processing functionality is not pre-loaded onto a device. Among their other uses, FPGA characteristics make them suitable for implementing cryptographic applications. However, while they offer this attack resistance and flexibility, they can also comprise yet another source of problematic issues, in that, in using the vendor's tools to configure them, a hardware engineer may introduce unintended vulnerabilities (fig 1).

SIMULATION

Typically, validation of such systems has relied primarily on simulation, a technique that can inform the hardware designer whether the system can indeed perform as intended, that is, whether the design's intent is preserved and reflected in its implementation and performance. As a validation tool, simulation, has its limitations, however. First, for stateful systems—that is systems with a large number of possible states—run times for simulations can become prohibitively long. For example, using a state-of-the-art simulator, the time required to simulate all possible combinations of inputs (each input can take a value of 0 or 1) for an n -bit adder is a matter of microseconds for $n = 4$, and still less than 100 milliseconds (one-tenth of a second) for $n = 8$. But the simulation run time jumps to over one hour for $n = 16$, and for $n = 32$ becomes a remarkable 585,000 years (fig. 2). The type of stateful systems involved in much of Sandia's mission space are often prohibitively difficult to simulate. And there is yet an even more crucial issue with simulation, namely that it can be somewhat blind to bugs, able to discover them only if clever test cases are used as the basis for guiding which system properties to simulate. There is no guarantee that disastrous bugs will all be discovered. While simulating a system can inform designers that it does what was intended in its design, it cannot necessarily also verify that it does not do what was not intended. In other words, simulation is generally not good at catching all unintended operational states—bugs—in a design.

In over 10 years working in the semiconductor industry for a leading FPGA company, Yalin Hu was repeatedly faced with this validation dilemma. As she poses it: "Through the years, I have seen customers facing the same question again and again: 'how confident can I be that what I want is what I get?'"




Figure 3. Simulation uses test cases to examine a system's state space in a somewhat incoherent fashion, and the probability of bug discovery is highly dependent on the cleverness of the test cases used to guide the simulations. Bugs (brown polygons) are easily missed. By contrast, formal verification systematically searches state space regions for bugs; eventually, the FV algorithm will search the entire state space, discovering all bugs.

Upon her arrival at Sandia, Yalin noted that there were critical systems caught in this limbo, whereby simulation was the primary tool being employed for verification, even while other validation and verification methodologies were beginning to enter the broader realm of FPGA-based system analysis. Using the vehicle of an early career LDRD research award, Yalin has been adopting and modifying other methods for verification of Sandia critical systems.

FORMAL VERIFICATION

These methods fall under the rubric of formal verification, a compendium of approaches employing rigorous mathematical proof that a hardware design satisfies certain specified properties. These include specific functionalities, timing properties, structural properties, and fault tolerance. Logically, this more-rigorous approach is necessary for mission-critical Sandia systems that are highly concurrent (simultaneously processing several data streams in parallel), that often operate in harsh environments (such as the high-radiation environment of outer space), and that, of course, are high-consequence systems, whose failure can have catastrophic outcomes, and whose fault tolerance is rather restricted. Formal verification is a technique to catch faults, that is, to verify that a system's




Figure 4. Comparison of runtimes for model checking of increasing RAM sizes, run as a single large problem (red) and as a set of decomposed problems (green). As the state space explodes with a problem of larger sizes, the runtime becomes prohibitive, but remains relatively stable for the decomposed problem.

array of reachable states does not include properties and states that were not intended in its design. It exhaustively explores all regions of a system's state space to uncover such incorrect system behaviors (fig. 3).

Under that formal verification umbrella are the techniques of model checking and theorem proving. Each has its advantages and its drawbacks. Model checking is a computationally intensive approach in that it entails an exhaustive examination of a system's collection of reachable states (its state space) to check that desired properties hold. This is, of course, a formidable task for complex (stateful) systems because of the large number of states that must be validated. Theorem proving is a method for logical derivation of system properties that is performed by mathematically defining the system's implementation. This is far from a trivial endeavor, requiring a significant investment of human intervention to construct such a mathematical definition of the system.

In order to facilitate the application of these more-rigorous methods to Sandia critical systems, this LDRD project is studying RAM (random access memory) a quite stateful system and is, more importantly, employing a decomposition approach, that is, breaking down a high-complexity problem to several lower-complexity problems.

DECOMPOSITION

Up to this point, there has been a relative dearth of research into this decomposition methodology, and therefore, Yalin Hu's LDRD research stands to provide Sandia critical systems with a great advantage in terms of their trusted functionality. The approach can be most easily expressed as decomposing a problem of size $P(N \times 2^M)$ into N problems of size $P(2^M)$. What this effectively does is to generate problems of a size that can be solved in a reasonable time period (fig. 4). Naturally, one must ultimately prove that such decomposition approaches have full validity, and part of this research is aimed at verifying that this is true. Such verification would pave the way for the use of the far more rigorous model checking with decomposition in validating critical Sandia systems.

The parallel formal verification approach of theorem proving is also a focus of this research, in that, currently, the theorems to mathematically represent a system are manually generated, consuming inordinate amounts of staff time. Were the decomposition approach successful in defining lower-complexity problems, a future goal would be to have the TP theorems be machine generated, and the hope is to develop algorithms to ultimately accomplish this. The possibility of automatizing the entire process of formal verification—model checking, theorem generation, and theorem proving—would represent an immense time savings to implement processes that would come much closer to ensuring the precise functionality of modern digital systems.

Finally, this research is not abandoning simulation, but is rather attempting to improve it. A process known as "assertion-based verification" can act as a bridge, introducing some features of formal verification into simulation to help devise simulation test cases in a more-rigorous fashion, thereby ultimately strengthening simulation as a verification tool.

It is rather obvious that within this critical systems domain, the consequences of getting it wrong, that is, of validating a system that may still, for example, possess a bug making possible a system state leading to control failure would be dire, potentially leading to loss of equipment, loss of life, and even catastrophic doomsday scenarios. These considerations boldly underline the critical nature of this research, work that should be strongly supported in every possible way. ■

For more information about the Sandia National
Laboratories LDRD Program, contact

Henry R. Westrich
LDRD Program Manager
(505) 844-9092
hrwestr@sandia.gov

Writer:
Vin LoPresti
Sandia National Laboratories

Graphic Design:
Chris Brigman
Sandia National Laboratories



Sandia National Laboratories is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND No. 2012-4334P

